Don’t Block My Stuff: Fostering Personal Object Awareness in Multi-user Mixed Reality Environments

TALHA KHAN, University of Pittsburgh, USA
DAVID LINDLBAUER, Carnegie Mellon University, USA

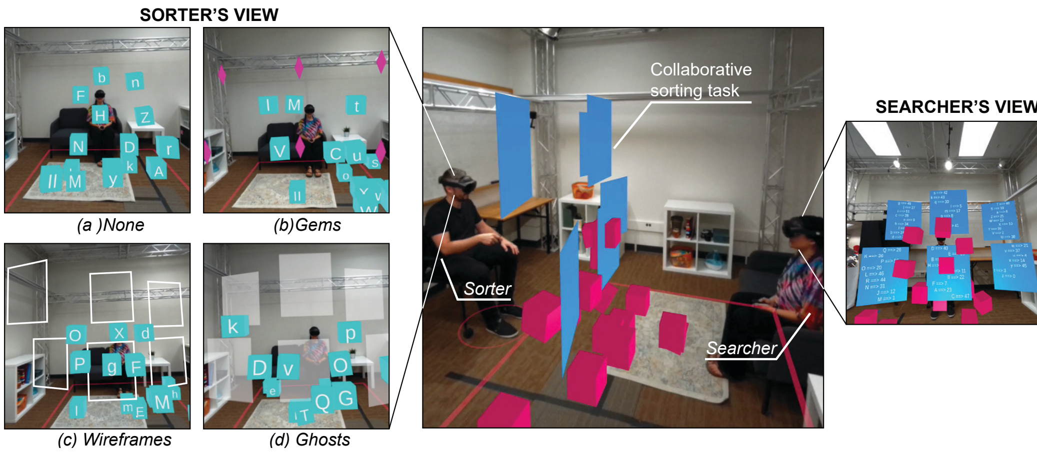

Fig. 1. We contribute an evaluation of different visualizations of personal interfaces in a collaborative Mixed Reality scenario. Two users perform a collaborative sorting task (*center*, from observer’s perspective). The *sorter* (left) sees cubes with letters that need to be sorted. The *searcher* (right) sees personal panels (blue) that reveal a mapping between the letters and the order of cubes, as well as the cubes but without the letters. While the *searcher’s* view of the panels and cubes remains the same, the *sorter’s* perspective of the panels changes with each visualization. The *sorter* experienced only one visualization at a time during the experiment. Note how the cubes placed by the *sorter* (right) blocks the *searcher* from seeing information on the panels. We investigate four different visualizations (a - d) of personal interfaces that enable the *sorter* to see where the panels of the *searcher* are, and their impact on occlusion and subjective ratings. *Image recorded live through Meta Quest Pro.*

In Mixed Reality (MR), users can collaborate efficiently by creating personalized layouts that incorporate both personal and shared virtual objects. Unlike in the real world, personal objects in MR are only visible to their owner. This makes them susceptible to occlusions from shared objects of other users, who remain unaware of their existence. Thus, achieving unobstructed layouts in collaborative MR settings requires knowledge of where others have placed their personal objects. In this paper, we assessed the effects of three visualizations, and a baseline without any visualization, on occlusions and user perceptions. Our study involved 16 dyads (N=32) who engaged in a series of collaborative sorting tasks. Results indicate that the choice of visualization significantly impacts both occlusion and perception, emphasizing the need for effective visualizations to

Authors’ Contact Information: Talha Khan, University of Pittsburgh, Pittsburgh, USA, muk21@pitt.edu; David Lindlbauer, Carnegie Mellon University, Pittsburgh, USA, davidlindlbauer@cmu.edu.

ACM 2573-0142/2024/12-ART526
https://doi.org/10.1145/3698126

enhance collaborative MR experiences. We conclude with design recommendations for multi-user MR systems to better accommodate both personal and shared interfaces simultaneously.

CCS Concepts: • **Human-centered computing → Visualization design and evaluation methods**; **Empirical studies in collaborative and social computing**; • **Security and privacy →** *Privacy protections.*

Additional Key Words and Phrases: Mixed Reality, Augmented Reality, Personal Interfaces, Visualization, Collaboration

ACM Reference Format:
Talha Khan and David Lindlbauer. 2024. Don't Block My Stuff: Fostering Personal Object Awareness in Multi-user Mixed Reality Environments. *Proc. ACM Hum.-Comput. Interact.* 8, ISS, Article 526 (December 2024), 24 pages. https://doi.org/10.1145/3698126

1 Introduction

Mixed reality (MR) systems present users with virtual content overlaid onto the physical world, often without restrictions on placement or size. In collaborative settings, MR facilitates the creation of shared interfaces. This enables efficient information sharing in a seamless manner [3] by providing shared visual context [14, 16]. Such interfaces allow for collaborative sense-making for a range of applications in domains such as healthcare [15], education [43], immersive analytics and information sharing [26, 34]. For instance, in a workplace setting, common examples of shared interfaces include slides and figures that a speaker shares with others during a presentation.

Besides shared interfaces in collaborative environments, MR also allows users to additionally and simultaneously have personal interfaces. These interfaces can contain sensitive information including messages, emails, calendar events, and more. Moreover, they can also contain task-specific information that enables individual users to make efficient contributions to the group, as in the case of asymmetric collaboration [38], where individuals have different roles and access to different sets of information. Displaying personal interfaces to everyone would lead to revealing private information, a consequence that is clearly undesirable [47].

Nevertheless, in many current MR systems, the focus is either on shared interfaces (visible to all) or on personal interfaces (visible only to the individual user). In settings where both these types of interfaces coexist, like a team an conducting exploratory data analysis task that incorporates shared and personal interfaces, *virtual-virtual conflicts* can arise, as illustrated in Figure 2. Users have no knowledge of where the personal interfaces of others are. If they move a shared interface, they can inadvertently position it in a way that occludes or overlap with a personal interface of another group member. This challenge is exacerbated in spatially-distributed MR layouts [10], where users can position personal interfaces such as ambient information displays (e.g., clock, weather app) or task-specific interfaces (e.g., dashboard) freely in space. To resolve such *virtual-virtual conflicts*, groups have to constantly re-position shared interfaces, or negotiate where elements can be placed, which can lead to frustration, conflict, and decreased performance.

We investigate solutions to mitigate *virtual-virtual conflicts* in MR environments that incorporate both shared and personal interfaces. We propose to use different visualizations to convey the presence of personal interfaces to other MR users in the space in a privacy-preserving manner. We explore the impact of three distinct visualizations, *gems*, *wireframes* and *ghosts*, and a baseline, *none*, without any visualization on *virtual-virtual conflicts*. All visualizations are shown in Figure 1. These visualizations were designed based on prior work, and convey different levels of information about personal interfaces to their non-owners. *Gems* indicate the existence and position of personal interfaces, but not their size or shape. *Wireframes* indicate the position, size and shape of interfaces, but omit the main content area so users can see through them. *Ghosts* indicate the position, size, shape of interfaces, and are mostly transparent with a slight level of opacity, limiting users from

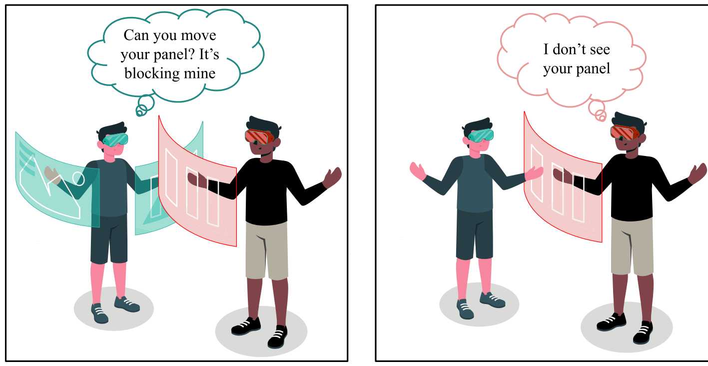

(a) (b)

Fig. 2. Example of a *virtual-virtual conflict*. Figure (a) illustrates the view of the person standing on the left, where both personal (shown in green) and shared (shown in red) interfaces are visible to them. Figure (b) depicts the view of the person standing on the right, note that the personal interfaces of the person on the left remain invisible to them.

fully seeing through them. The baseline without visualization represents current MR systems, where by default personal interfaces are only visible to their owner. For all visualizations, the content of personal interfaces is never shared to preserve privacy.

The core contribution of our work is an investigation into the impact of individual visualizations on *virtual-virtual conflicts* in collaborative scenarios, and to gain insight into users' perceptions and preferences regarding these visualizations.

To achieve this, we conducted a study with 16 dyads (N=32) during which participants performed an asymmetric collaborative task. Pairs of participants were tasked with sorting a series of cubes, illustrated in Figure 1. One participant, the *sorter*, arranged these cubes (the shared interfaces), while the other, the *searcher*, determined their correct order using information from virtual panels (the personal interfaces). In our experiment, the cubes represent shared interfaces that are visible to both the *sorter* and the *searcher*; and the panels are personal interfaces, visible only to the *searcher*, and represented by aforementioned visualizations.

Our results revealed notable differences in visualizations for reducing occlusions. The *wireframes* and *ghosts* visualizations proved to be effective in reducing occlusions, while the *gems* visualization occasionally exacerbated the issue. Feedback collected from subjective questionnaires and interviews also supported these findings. *Ghosts* were preferred by participants, followed by *wireframes*, compared to *gems* and *none*, albeit with nuances in individual preferences. The *none* condition, for example, though less effective in mitigating occlusions, was favored by a few *sorters*, again underscoring the diverse preferences within MR users.

Based on our findings, we distill a set of design guidelines. We highlight that *ghosts* and *wireframes* mitigate occlusion compared to simpler visualizations such as *gems*, or the conventional *none* condition. We believe our study is a step towards making collaboration in MR more efficient, and enabling a wide range of applications in which users interact with both shared and personal interfaces simultaneously.

2 Related Work

2.1 Overview of Mixed Reality Systems

2.1.1 Collaborative Mixed Reality Systems. MR technology has seen widespread use in collaborative systems and environments, enhancing user collaboration capabilities [12, 48]. Research has primarily focused on integrating MR with traditional displays like tablets and large wall displays to improve interaction and data analysis. Examples include MARVIS [26], which extends tablets with MR for data interaction, and ART [5], which merges AR with touch-interactive tabletops for analyzing multidimensional data. Dataspace [7] combines large displays with Augmented Reality (AR) and Virtual Reality (VR), providing various interaction modes (e.g., touch, gaze, gesture) to create flexible, immersive analytics environments.

In addition to enhancing traditional displays, standalone MR systems also facilitate collaborative work. Mahmood et al. [34] introduced a remote MR system for geospatial data analysis that enables more efficient communication than conventional displays. FIESTA [45] explores VR-based collaborative analytics in immersive environments. MR's scope extends to content creation and design, as exemplified by TransVision [46] and C-Space [49], which support collaborative 3D design activities. Radu et al. [42] introduced a novel MR system that facilitates remote collaboration in physical maker space activities, leveraging AR, VR, and 3D scanning technologies. The educational potential of MR is still being explored [22], with systems aiming to optimize learning experiences while addressing limitations like tunnel vision [43]. In healthcare, the ARTEMIS platform enables expert surgeons to mentor novices via AR/VR [15].

To further enhance collaborative experiences in MR, researchers have integrated interactive mechanisms like avatars, gaze, and gestures. For instance, providing consistent gaze, gesture, and pointing cues through an adaptive avatar (e.g., transforms size and orientation) enhances social presence and reduces task load for remote collaboration between AR and VR users [39]. Including eye gaze and hand gestures has been shown to increase co-presence and joint attention in both co-located and remote collaborative settings [1, 21], while virtual objects used as spatial cues have been effective in enhancing grounding and reducing perceived workload [37]. Our work contributes by enabling groups of co-located MR users to collaborate while avoiding *virtual-virtual* conflicts in environments where shared and personal interface elements co-exist.

2.1.2 Personal Mixed Reality Systems. Beyond its collaborative applications, MR technology has been extensively explored for individual use. The applications of these individual systems often mirror their collaborative counterparts. For example, there are MR systems designed for immersive analytics [19], academic endeavors like paper reading [28], design processes [13], healthcare [23, 24], and more. We refer readers to the surveys by Kim et al. [25] and Zollman et al. [52] for discussions on various application domains.

A notable challenge for personal MR systems is the significant effort that users have to invest in maintaining the sensible layout of personal interfaces [8]. To address this, researchers have proposed automated techniques for view and layout management. Bell et al. [2], for instance, introduced methods to manage the layout of objects in 3D virtual and augmented environments to maintain spatial (e.g., distance between objects) and visibility (e.g., A should not occlude B) constraints provided by users. Furthermore, multi-modal data, such as speech, has been employed to adjust the position and opacity of virtual interfaces based on conversation topics [11]. Factors like cognitive load, task type, and the user's environment have been jointly considered to optimize MR user interfaces in a context-aware manner [8, 29].

These works highlight the need for efficient techniques that enable users to constantly view relevant content while avoiding clutter and occlusions. In collaborative settings, however, these factors are not yet well explored. It is yet unclear how much of others' virtual elements users would like to see, and how much they are willing to share. This information is necessary for building any kind of advanced adaptive collaborative MR system. Although there have been discussions and guidelines proposed for such integrations from the field of cross-device collaboration [4], actual system implementations and empirical testing in collaborative contexts have yet to gain traction.

2.2 Visualizations in Mixed Reality

In MR, visualization techniques have been utilized for various purposes such as enhancing depth perception with visual cues and reducing visual clutter through abstractions.

Livingston et al. [31] explored different visualization techniques in MR and assessed their impact on enhancing depth perception in case of multiple physical and virtual objects. The techniques included varying drawing styles – wireframe, solid fill, and a combination of both.

In the domain of security and privacy, Ruth et al. [47] introduced the concept of ghost objects, semi-transparent AR objects that signal the presence of a virtual object to other users in a shared space without disclosing detailed information. This technique effectively balances privacy with the awareness of other users' activities in a collaborative environment. Similarly, Rajaram et al. [44] proposed coarse-grained visualizations to manage the trade-off between information awareness and privacy in multi-user AR settings. These visualizations simplify the details of virtual objects, preserving confidentiality while still allowing users to recognize their presence and shape.

In the area of Diminished Reality, Taylor et al. [50] explored the use of outline visualizations superimposed on physical objects (e.g., a robot arm). These visualizations render the physical object semi-transparent but retain an outline, enhancing users' spatial awareness and tracking of the object's position and movements without obscuring the background scene. Cheng et al. [9] also investigated various visualization techniques within the context of diminished reality to address information overload. They tested different approaches such as blurring, reducing opacity, and outlining to selectively omit physical objects from the user's view, thus reducing clutter in the environment and allowing users to focus attention on relevant tasks. Our work is inspired by the visual manipulations of these works to indicate the presence of personal virtual elements while varying the degree of information we provide to users.

2.3 Contextualizing Our Work

While MR has been extensively studied in collaborative settings, most research focuses on fully symmetric environments where all users see all virtual elements. Few works have explored environments in which both shared and personal interfaces coexist. The sharing of personal interfaces also presents conflicting views. On the one hand, Ruth et al. [47] suggest that personal views should be shared with other MR users in a privacy-preserving manner in order to socially signal to other users that they are busy interacting with private content. On the other hand, personal views are oftentimes seen as a main building block of MR, allowing users to interact with their private data without others knowing [45].

Despite the importance of personal interfaces, there is a clear need to foster their awareness. Previous works [27, 40, 41] have noted that users frequently place virtual objects within each other's personal space [17], resulting in discomfort [32, 51]. However, none of previous works have investigated the use of visualizations in enhancing awareness or mitigating *virtual-virtual conflicts*, where users accidentally interfere with each others' personal virtual content.

The work most similar to ours is by Jackson et al. [20], who investigated techniques for communicating personal workspace boundaries between co-located AR users performing independent, non-collaborative tasks. They evaluated three different workspace guardian visualizations: showing the full virtual content, displaying just the bounding box outlines of the content, and allowing users to self-define a boundary region. Their study measured the impact of these techniques on personal workspace encroachments (i.e., conflicts) by other users. However, in their work they focus on conflicts between a user and virtual interfaces, as opposed to our work, which focuses on conflicts among virtual interfaces only.

3 Method

We conducted an empirical study to investigate the influence of different visualizations of personal interface elements in a collaborative MR setting. Specifically, we were interested in how much information about others' interface elements users would like to see, which visualizations prevented occlusion, and how they influenced users' behavior and preferences. Pairs of participants, each having partial awareness about the task, were instructed to perform a collaborative sorting task. Participants were only instructed about the goals of the task and basic interactions. They were not informed about the potential occurrence of occlusions or any other aspects related to collaboration.

3.1 Design

Our study utilized a within-subjects design with a single independent variable VISUALIZATIONS, with four levels: *none*, *gems*, *wireframe*, and *ghosts*. The visualizations are shown in Figure 1, and Figure 3.

Pairs of participants performed the collaborative sorting task once with each VISUALIZATION. During the task, we distinguished between two roles. The *searcher* could search for information that was essential for completing the task. This information was only visible to them on their personal objects. The *sorter* was tasked to follow the instructions of the *searcher* and sort objects in the space accordingly. While the objects of the *sorter* were shared and visible to the *searcher*, the personal objects of the *searcher* were presented with different visualizations to the *sorter*.

3.1.1 Visualizations. We chose the VISUALIZATIONS based on multiple criteria: the degree of visual interruption they cause, the amount of spatial information (such as position, size, and shape) they reveal, and their effectiveness in preserving privacy. Figure 3 summarizes the properties of the VISUALIZATIONS used in our experiments. The purpose of these visualizations was to convey existence of personal virtual objects to others within the space, presenting this information at different levels of detail.

- ***None***: Personal objects of the *searcher* remained invisible to the *sorter*. This condition acted as our baseline, mirroring the *status quo* of numerous MR systems where personal objects are visible only to their owner and hidden from other users.
- ***Gems***: Personal objects were displayed to the *sorter* as gems (small diamond-shaped 3D objects). The gems were positioned at the center of the *searcher's* objects and were roughly the size of a cuboid with length 10 cm, height 20 cm, and depth 10 cm. This visualization conveys the presence and position of personal objects, while simultaneously minimizing

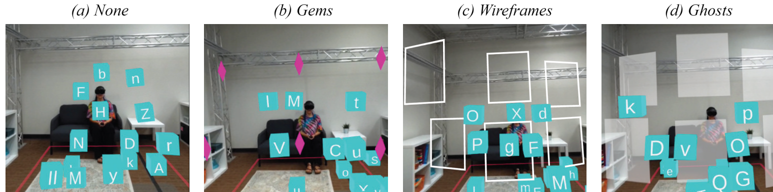

	(a) None	(b) Gems	(c) Wireframes	(d) Ghosts
Reveals **size** of personal interfaces	×	×	✓	✓
Reveals **position** of personal interfaces	×	✓	✓	✓
Reveals **shape** of personal interfaces	×	×	✓	✓
Privacy Preservation	High	Medium	Medium/Low	Medium/Low
Visual Disruption	None	Low	Medium	High

Fig. 3. Table showing all the VISUALIZATIONS and their corresponding properties.

visual disruption and safeguarding privacy by concealing the shape and size of the personal virtual objects.

- ***Wireframes*:** The *sorter* was only presented with the *wireframes* (essentially the outline) of the *searcher's* personal objects. This visualization conveys the position, shape, and size of personal interfaces, without significantly obstructing the user's view of the surrounding environment, thereby causing a moderate level of visual disruption, albeit at the expense of reduced privacy. Wireframes are inspired by prior work on occlusion [31] and Diminished Reality [9, 50].
- ***Ghosts*:** Personal objects were presented to the *sorter* as semi-transparent (translucent) objects. This approach effectively communicated the location, shape, and size of the virtual items, albeit with a high degree of visual disruption and compromised privacy. This visualization served as a reminder to the *sorter* that the interface of the *searcher* was not hollow. Ghosts are inspired by prior work on occlusion [31], Diminished Reality [9, 50] and AR privacy [47].

In all of the visualizations, no personal information was visible to the other user.

3.2 Task

Our experimental task required participants to collaboratively solve a sorting assignment. The task involved two primary types of virtual objects: cubes and panels. Participants were asked to sort 15 cubes based on a task-specific scheme. Each cube had a lowercase or uppercase letter displayed on it, randomly chosen for each experimental condition from the full set of 52 letters (a-z, A-Z). The panels displayed associations linking the letters to numbers e.g., *c = 15*. Six panels in total provided associations for all 52 letters. Participants were asked to sort the cubes in ascending order across three rows, with each row holding five cubes. The order of a cube was determined by the number corresponding to the letter displayed on it. These cubes behaved like physical objects in terms of occlusion, i.e., placing a cube in front of a panel would occlude the panel, and vice versa.

3.2.1 User Roles. Our task was asymmetric in nature with users taking on different roles and each having partial awareness of the task. The users were seated face-to-face to each other, a common *f-formation* configuration for collaboration [35, 36], while performing the task. Users were approximately 2 m apart, giving each user roughly 1 m of personal space, as shown in Figure 1. The *sorter* was responsible for interacting with the cubes i.e., moving them and arranging them in the correct order, through standard distant interaction (i.e., raycasting to select; controller-grab to manipulate) using the headset controller. The *searcher* was responsible for searching information relevant to letter-to-number associations on panels displayed in front of them, shown in Figure 1. The *searcher* was not able to interact with the cubes. Information on the panels was randomly shuffled every seven seconds. This required the *searcher's* consistent engagement and search effort during the task, inducing a slightly higher cognitive load, especially when dealing with potential interference between virtual objects.

While the cubes were visible to both the *sorter* and the *searcher*, the letter printed on them could only be seen by the *sorter*. The panels were always visible to the *searcher*, however, the *sorter* saw different visualizations of the panels under different experimental conditions. The *sorter* was informed about the presence of the panels even in instances where they were invisible (i.e., the *none* condition). We chose this design, since we assumed a scenario where others are not naive about the existence of virtual elements; they just do not have specific information (e.g., content, position, size). This is, for example, similar to situations where users are aware that others interact with their smartphone, but do not know what they are doing specifically. MR exacerbates this challenge, since users no longer have information about position and size of the interface too.

3.2.2 Task Choice. We selected this task because it aligns well with our research objectives. It required collaboration between users who each have only partial solution knowledge — specifically, the letters on shared objects (i.e., cubes) and the letter-to-number associations on personal objects (i.e., panels) were visible to only one user at a time. This design ensured that occlusions on the panels impacted both users, thereby evoking more natural reactions. Additionally, the task enabled enabled the application of different VISUALIZATIONS to the virtual interfaces (i.e., panels). This task also ensured that *virtual-virtual conflicts*, the main focus of this work, would indeed occur, allowing us to study them in a controlled manner. We chose a face-to-face formation, since it is similar to current co-located collaboration scenarios where users gather around a conference table or sit in an open space area. We hope to explore other configurations such as side-by-side setups (e.g., similar to users working on a shared whiteboard) in the future.

3.2.3 Experimental Layout. The cubes had a side length of 0.17 m, while the panels had a side length of 0.5 m and a height of 0.6 m. The panels were 2D windows, typical for many MR applications. The panels were designed to ensure they could accommodate all task-relevant information in a legible way. The size of the panels was kept constant across conditions to allow us to quantify occlusion in a comparable manner, as shown in Figure 4. The layout of the panels was similar to common productivity layouts for MR [30, 33].

At the beginning of each experimental trial, the cubes were randomly spawned at a vertical distance between 0.2 m and 0.3 m from the ground. This ensured that all conditions started without occlusions. The panels were set up in two rows: the bottom row's center was 0.75 m vertically above the ground, while the top row's center was at 1.5 m. Each panel was assigned a depth from the *searcher's* position, either 1 m (near) or 1.2 m (far), with an equal number of panels (three each) at both depths. These depths were roughly within the personal space of the *searcher* (cf. Hecht et al. [18]). Horizontally, the panels extended a range from −0.8 m to 0.8 m. After setting the initial panel positions, a random adjustment of ± 0.1 m was applied to their vertical, horizontal, and depth positions. Finally, the panels at each end were rotated by 15 degrees to form a semi-circular arrangement facing the *searcher*.

3.2.4 Other Virtual Objects. We incorporated a subtle square-shaped virtual ground marker of side length 2 m, depicted in red in Figure 1. The center of this marker was positioned roughly 1 m in front of both users and acted as a soft boundary. The *sorter* was asked to keep the cubes roughly within the limits of the marker while performing the task. The marker ensured that *sorters* positioned the cubes directly in front of them, which was crucial for facilitating the face-to-face formation for the collaborative scenario our study intended to investigate.

Additionally, there was a seventh panel, only visible to the *searcher*. This additional panel informed the *searcher* how their virtual objects (i.e., the panels) were being displayed to the *sorter*. We chose to show this additional panel, because in a real-world setting, a user would know how their personal interfaces were being presented to others. Also, since the view of the *searcher* remained consistent throughout the experiment (as shown in Figure 1), this additional panel helped them differentiate among experimental conditions when ranking them.

3.2.5 Task Order. All participants experienced a predetermined sequence for the experimental conditions, starting from VISUALIZATIONS that revealed less details about the personal interfaces (i.e., high privacy preservation), and progressing to VISUALIZATIONS that exposed more (i.e., low privacy preservation). Participants always started with *none*, followed by *gems*, then *wireframes*, and finally *ghosts*. We chose this methodology to prevent prematurely sharing details about the personal interfaces (e.g., their shape or size) with participants, which effectively would have made them *oracles*. Each visualization incrementally contained more information about personal objects than the previous one: *none* revealed no information about the personal objects; *gems* revealed the position; *wireframes and ghosts* additionally revealed size and shape, with *ghosts* providing slightly more insight into shape than *wireframes*, see Figure 3.

We extensively piloted different variations of the experimental design. During the pilots, when participants were presented with visualizations that provided more insights first, for example first *ghosts* then *gems*, they altered their behavior to incorporate the "hidden" information in the latter condition: they preemptively placed cubes on the ground or ceiling during *gems*, since they could anticipate the rough layout of the panels from the previous condition, but did not have enough information to precisely avoid occlusions. They did so even though it led to physical discomfort (e.g., participants reported neck strain from looking at the ceiling). We therefore do not consider the "oracle-like" behavior as realistic.

Finally, we chose a within-subjects design over a between-subject design to gather qualitative feedback on how the different visualizations compared to each other. We believe that our current findings would be enriched by replicating our study as a longitudinal between-subject experiment.

Mitigating Order Effects. To mitigate order effects, we introduced random variations to the positions of the panels (see Section 3.2.3). For instance, in one trial, the bottom middle panel might be positioned at a depth of 0.9 m from the *searcher's* position while the top middle panel could be at 1.3 m, creating a depth variance of 0.4 m. Using this strategy, replicating the layout from previous conditions — where cubes did not obstruct the panels — could result in occlusions in the current condition. This made it necessary for participants to generate a new arrangement for the cubes and in turn effectively mitigates the order effects. However, we kept the general layout similar to productivity environments, utilizing a majority of available space in the frontal visual field. We hope to expand our experiment to other types of layouts (e.g., casual layout for a collaborative board game) in the future.

3.3 Apparatus

Our platform was implemented using Unity 2021.3.24f1 and the Oculus Integration SDK v53.0. We implemented a Python server for communication between the headsets. This server ran on a Dell Alienware x17 R1 computer (equipped with an 11th Gen Intel(R) Core(TM) i9-11980HK processor with 64Gb of RAM) and was responsible for the initiation of the different experimental conditions and data collection. All apparatus was connected to the same local network.

Two Meta Quest Pro headsets were used by participants during the experimental process. Hand interactions were disabled for both headsets and interactions with virtual objects were only possible through the headset controllers. A third Meta Quest Pro headset was used by the experimenter to observe the participants in order to ensure that they were adhering to the study guidelines. All headsets operated in passthrough MR mode, superimposing virtual objects onto the camera feed of the real world, mimicking an AR-like experience. A live recording of the view is shown in Figure 1. All headsets were calibrated for each participant to share the same virtual space to synchronize the placement of virtual elements.

3.4 Participants

We recruited 32 paid participants (21 male, 11 female), resulting in a total of 16 dyads, from a local university and an external recruiting platform. Participants' ages ranged from 18 to 40 years ($M = 24.6$, $SD = 5.08$). Participants had a median experience with AR of 2 ($M = 1.87$, $SD = 0.94$), and VR of 2 ($M = 2.15$, $SD = 1.08$), both on a scale from 1 (no experience) to 5 (expert). According to self-reports, all participants had either normal or corrected-to-normal vision, with 19 of them wearing glasses or contacts.

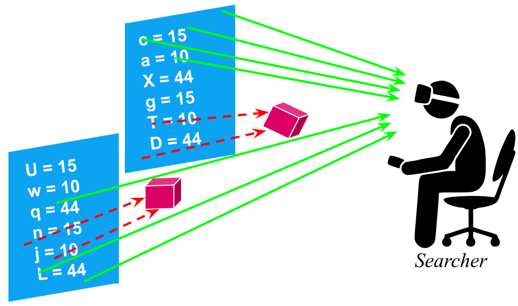

Fig. 4. Figure illustrating how occlusion is defined for a panel. The red rays are blocked by cubes before reaching the *searcher*, therefore are marked as occluded. The green rays directly reach the *sorter* and are considered unoccluded. Occlusion was measured by finding the ratio of red rays to all (red + green) rays. For each panel, we emit 100 equally distributed rays toward the *searcher*'s head.

3.5 Procedure

The study was conducted inside a quiet experimental space. After signing the consent form and completing the demographics survey, participants were randomly assigned roles of *searcher* or *sorter*. We then initiated a training task to familiarize participants with the interaction mechanics, such as grabbing and moving a virtual object using the controller. Afterwards, we introduced the first of four experimental conditions. Participants undertook the task at their own pace and alerted the researcher upon completion. The researcher checked the participants' answers and informed them whether they were correct or incorrect. The task was considered completed once the correct order was achieved. After every condition, participants removed their headsets to complete a questionnaire on a 5-point Likert-type scale (1 strongly disagree to 5 strongly agree; 1 very low to 5 very high) addressing aspects like occlusion, distraction, comfort, and collaboration. To conclude the study, we asked both participants to rank the different conditions individually. Additionally, we held a semi-structured interview to understand the reasons behind their rankings and to gather general feedback on the study.

The study was approved by the local university's Institutional Review Board (IRB). The study spanned approximately one hour. Participants were compensated with an honorarium of $25 (USD).

3.6 Data Collection

All data was timestamped and collected twice every second on the computer running the server. We collected data containing the position and orientation of both users, virtual object attributes such as their position, orientation, and scale, amount of occlusion of each panel, specific events like the start or end of an experimental condition, and object interactions (e.g., object x grabbed/released).

We calculated the amount of occlusion for each timestamp by emitting 100 rays from each panel (uniformly distributed) towards the *searcher*, illustrated in Figure 4. If a ray intersected with a cube before reaching the *searcher*, it was considered occluded. The amount of occlusion is the ratio of occluded to total rays. We analyzed the occlusions present in the six panels that illustrated associations between letters and numbers. Panels could only be occluded by cubes that the *sorter* moved. The seventh panel, which informed the *searcher* about the current condition, was excluded from this analysis.

4 Results

All statistical analyses were conducted using SPSS (version 28.0.1.1). If the data satisfied the assumptions of normality and sphericity, ANOVA was applied. Dunn tests with Bonferroni correction were used for post-hoc analysis. For posthoc tests, we only report statistically significant pairwise differences. In cases where these assumptions were not met, the Friedman test was used. We transcribed all interviews using an automated transcription service and made manual corrections as needed to rectify transcription errors.

4.1 Occlusion

In the following, we quantify occlusion with different metrics to provide a holistic view, capturing the number of panels occluded, the extent of the occlusion, and the duration of the occlusion.

4.1.1 Occlusion Duration Threshold. We first calculated a threshold for the duration a panel needed to be occluded by a cube to be counted as occluded for our calculations, denoted as $T_{\text{threshold}}$. We introduced this threshold to distinguish short fleeting occlusions from more genuine substantial ones that would negatively impact user experience. Short occlusions typically occurred while the *sorter* was moving a cube to its position. In contrast, more extended occlusions occurred when the cube was positioned to block the panel for an extended period. Leveraging the threshold results in a more meaningful calculation for metrics such as number of occluded panels, amount of occlusion, and duration of occlusion.

For each session, we first identified the maximum duration, $T_{\max}$, a cube was grabbed by users. $T_{\text{threshold}}$ was calculated by taking the average of all $T_{\max}$ durations across all sessions. The value of $T_{\text{threshold}}$ was 10.39 seconds, i.e., the average maximum grab time. A panel was counted as occluded if there was a consecutive time period of occlusion during a session that exceeded $T_{\text{threshold}}$.

4.1.2 Number of Occluded Panels. We measured the number of occluded panels for each condition across all sessions, shown in Figure 5a, i.e., the number of panels that were considered occluded based on the threshold criteria. We observed that with more informative visualizations the number of occluded panels decreased. A Friedman test revealed a significant main effect for VISUALIZATIONS on the number of occluded panels ($\chi^2(3) = 28.579$, $p < 0.001^{***}$). The coefficient of concordance, Kendall's W, indicated a good agreement [6] ($W = 0.595$). Posthoc tests indicated that *none* led to significantly more occlusions than *wireframes* ($Z = 1.469$, $p = 0.0077^{**}$) and *ghosts* ($Z = 1.719$,

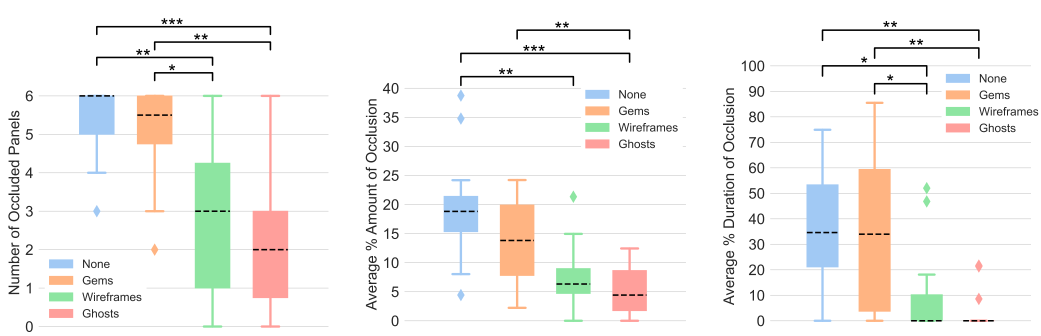

(a) Number of occluded panels (b) Amount of occlusion (c) Duration of occlusion

Fig. 5. Results of occlusion across conditions. Horizontal bars indicate pairwise statistically significant differences when main effects were present.

$p < 0.001^{***}$). Similarly, *gems* led to more occlusions than *wireframes* ($Z = 1.406$, $p = 0.0123^{*}$) and *ghosts* ($Z = 1.656$, $p = 0.0017^{**}$).

4.1.3 Amount of Occlusion. For each session, we calculated the average percentage amount of occlusion. We only considered panels in this calculation if they exceeded the temporal occlusion threshold, $T_{\text{threshold}}$ (see Section 4.1.1). For each of these panels, we tracked the number of occluded rays at every timestep (Section 3.6). The amount of occlusion per panel is the average number of occluded rays over time. The occlusion per session is the average of the per-panel occlusion. We decided to include occluded panels based on the threshold to avoid artificially deflating the average occlusion amount (i.e., panels without occlusion would have decreased the average in a misleading way).

The results for the amount of occlusion are shown in Figure 5b. With more informative visualizations, we observed a reduction in the average percentage amount of occlusion. A one-way repeated measures ANOVA indicated a significant effect for VISUALIZATIONS on the amount of occlusion, $F_{(3,45)} = 16.917$, $p < 0.001^{***}$, $\eta_p^2 = 0.53$. The large effect size suggests that VISUALIZATIONS had a substantial influence on the amount of occlusion. In a subsequent post-hoc analysis, we found that the *none* condition led to significantly more occlusion than *wireframes* (μ difference = 11.746, $p = 0.0012^{**}$) and *ghosts* (μ difference = 13.951, $p < 0.001^{***}$). *Ghosts* exhibited less occlusion than *gems* (μ difference = 8.555, $p = 0.0059^{**}$).

4.1.4 Duration of Occlusion. For each session, we computed the average percentage of time a panel was occluded. We only considered panels in this calculation if they exceeded the temporal occlusion threshold, $T_{\text{threshold}}$. The duration of occlusion is the average time each panel was considered occluded during a session. The average duration for the session was the average per-panel occlusion duration.

The results for the duration of occlusion are shown in Figure 5c. A Friedman test highlighted a significant effect for VISUALIZATIONS on average percentage duration of occlusion ($\chi^2(3) = 26.368$, $p < 0.001^{***}$). Kendall's W showcased a good agreement ($W = 0.549$). In the posthoc tests, *none* exhibited longer occlusion duration than *wireframes* ($Z = 1.344$, $p = 0.0194^{*}$) and *ghosts* ($Z = 1.594$, $p = 0.0028^{**}$). *Gems* led to longer occlusions compared to *wireframes* ($Z = 1.406$, $p = 0.0123^{*}$), and *ghosts* ($Z = 1.656$, $p = 0.0017^{**}$).

4.1.5 Proxy Measures. During our experiments, various interactions between the *searcher* and the *sorter* emerged. Commonly, the *searcher* would ask the *sorter* to adjust the cubes' position when their view was blocked, while the *sorter* periodically inquired about any occluded panels. Occasionally, the *searcher* adjusted their viewpoint to glimpse the blocked information from an alternative angle. For more informative visualizations, both mentions of occlusions and head movements due to blockages were reduced.

Time Away from Modal Position. To view occluded information on the panels, *searchers* had to adjust their head position and move away from their modal position, i.e., the position where they most frequently sat during each experimental condition. We thus analyze how much time away from this modal position *searchers* spent as a proxy metric for occlusion.

We calculated this metric using *k-means* clustering and took the centroid of the cluster with the most points as the modal position. In all sessions, we found the optimal number of clusters was $k = 3$, based on the elbow method. We set a 10 cm threshold to identify head movements caused by occlusions. Whenever the *searcher's* position strayed more than 10 cm from the modal position, we counted that timestep towards our calculation. Figure 6a displays the distribution of the percentage of time the *searcher* spent outside this modal position. Conditions with more informative visualizations resulted in the *searcher* making fewer occlusion-related movements.

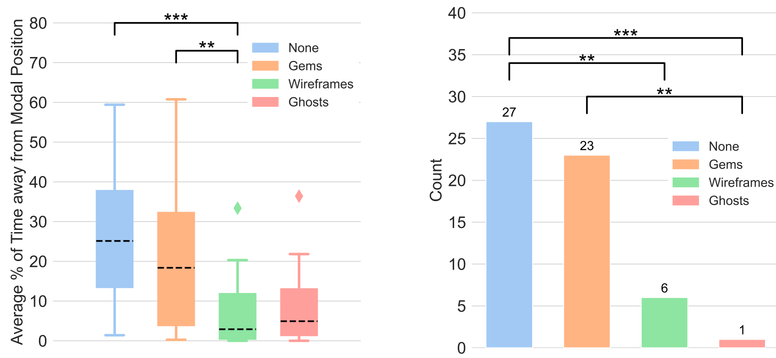

(a) Time spent away from modal position

(b) Number of verbal occlusion references

Fig. 6. Plots for proxy measurements for occlusion

While the average time away from the modal position was relatively high for *none* ($M = 26.9\%$, $IQR = 24.461$) and gems ($M = 21.4\%$, $IQR = 26.521$), the number converged to below 10% for *wireframes* ($M = 7.53\%$, $IQR = 11.569$) and *ghosts* ($M = 8.81\%$, $IQR = 11.849$). A Friedman test highlighted a significant effect for VISUALIZATIONS on the average percentage of time spent away from modal position ($\chi^2(3) = 18.698$, $p < 0.001^{***}$, $W = 0.390$). For post-hoc analysis, we conducted a series of Dunn tests with Bonferroni correction. We found that participants spent significantly more time away from the modal position in the *none* ($Z = 1.719$, $p < 0.001^{***}$) and *gems* ($Z = 1.531$, $p = 0.0047^{**}$) condition compared to *wireframes*. These findings go in line with an increase in occlusion between *none* and *gems* compared to *wireframes*.

Verbal Reference of Occlusion. We reviewed the transcribed audio recordings from the sessions and counted how often users discussed (verbally referenced) panel occlusions, for example, (*"The three in the middle top row are blocking; just move them a little bit lower"*, P2-*searcher*). Figure 6b displays the total references for each condition. Each condition occurred 16 times during our experiment, once per dyad. The *none* condition had the highest number of occlusion references, followed by *gems*, *wireframes*, and *ghosts*. Every session contained at least one occlusion reference. While the average number of times a verbal reference was made related to occlusion was low across all conditions, *none* ($M = 1.69$, $IQR = 0.85$), *gems* ($M = 1.44$, $IQR = 1.17$), *wireframes* ($M = 0.38$, $IQR = 0.48$), and *ghosts* ($M = 0.06$, $IQR = 0.24$), a Friedman test highlighted a significant effect for VISUALIZATIONS on number of verbal occlusion references ($\chi^2(3) = 32.712$, $p < 0.001^{***}$, $W = 0.681$). For post-hoc analysis, we conducted a series of Dunn tests with Bonferroni correction. We found that participants made significantly more verbal occlusion references in the *none* ($Z = 2.031$, $p < 0.001^{***}$) and *gems* ($Z = 1.531$, $p = 0.0047^{**}$) condition compared to *ghosts*, and *none* ($Z = 1.594$, $p = 0.0028^{**}$) condition compared to *wireframes*.

4.2 Task Performance

We report the time, measured in seconds, taken by participants to complete the task. The task had no time limit and was considered complete only when participants achieved the correct sorting order. Participants used a variety of strategies for sorting, with some groups being more efficient than others. For the *none* condition, $M = 777.84$, $IQR = 268.26$; for the *gems* condition, $M = 579.96$, $IQR = 243.40$; for the *wireframes* condition, $M = 518.28$, $IQR = 201.99$; and for the *ghosts* condition,

$M = 426.969$, $IQR = 185.88$. It is important to note, however, that the task completion time was not the main focus of our study; rather, the goal was to provide a comprehensive analysis of the impact of VISUALIZATIONS on occlusions and the resulting impact on user's perceptions.

4.3 Subjective Ratings

After each experimental condition, we asked participants to complete a survey about their perception. Given the asymmetric nature of our task, we analyzed responses from the *searcher* and the *sorter* separately. Figure 7 displays the survey responses of the *searchers*, and Figure 8 the responses of sorter. We carried out Friedman tests to assess statistical significance.

4.3.1 Searcher's Analysis. In summary, the *ghosts* condition was perceived as less distracting, less interfering, more comfortable, and conducive to collaboration than *none* and *gems*. We found significant main effects across several dimensions of the *searcher's* survey responses. We found a significant main effect for distraction ($\chi^2(3) = 20.284$, $p < 0.001^{***}$, $W = 0.432$). Participants perceived *ghosts* as significantly less distracting than *none* ($Z = 1.496$, $p = 0.0077^{**}$) and *gems* ($Z = 1.496$, $p = 0.0077^{**}$). We found a significant main effect for interference ($\chi^2(3) = 22.907$, $p < 0.001^{***}$, $W = 0.477$). Again, ghosts were perceived as resulting in less interference than *none* ($Z = 1.344$, $p = 0.0194^{*}$) and *gems* ($Z = 1.406$, $p = 0.0123^{*}$). The responses to questions *'Comfort with my content positioning'* and *'Comfort with other's content positioning'* exhibited significant main effects ($\chi^2(3) = 11.182$, $p = 0.0107^{*}$, $W = 0.233$ and $\chi^2(3) = 11.244$, $p = 0.0104^{*}$, $W = 0.234$); the former did not show any significant post-hoc differences, the latter revealed significant differences for the *none* and *ghosts* condition with ($Z = -1.250$ and $p = 0.0370^{*}$). For *'Effect of other's content on collaboration'*, there was a significant main effect ($\chi^2(3) = 19.697$, $p < 0.001^{***}$, $W = 0.410$). Post-hoc tests showed that *ghosts* was ranked significantly higher than *none* ($Z = 1.564$, $p = 0.0037^{**}$) and *gems* ($Z = 1.406$, $p = 0.0123^{*}$). Lastly, despite the significant main effect for *'Irritation from other's virtual content'* with ($\chi^2(3) = 14.607$, $p = 0.021^{**}$, $W = 0.304$), post-hoc comparisons did not indicate pair-wise difference.

4.3.2 Sorter's Analysis. In summary, the *ghosts* and *wireframes* conditions showed differences in awareness compared to the *none* condition. However, despite numerous significant main effects in other categories like distraction, blocking extent, and irritation, post-hoc tests often did not indicate pair-wise differences between specific conditions.

Results revealed a significant main effect for distraction ($\chi^2(3) = 14.341$, $p = 0.0024^{**}$, $W = 0.299$). Post-hoc comparisons did not reveal any significant differences. We found a significant main effect for the question *'Awareness of other's virtual content'* ($\chi^2(3) = 33.206$, $p < 0.001^{***}$, $W = 0.692$). Post-hoc analyses further highlighted significant differences between the *none* and *ghosts* conditions with ($Z = -2.062$, $p < 0.001^{***}$) and between the *none* and *wireframes* conditions with ($Z = -1.937$, $p < 0.001^{***}$).

Results further revealed main effects for the questions *'Extent of blocking other's content'* ($\chi^2(3) = 11.531$, $p = 0.0091^{**}$, $W = 0.240$), and *'Irritation from other's virtual content'* ($\chi^2(3) = 10.186$, $p = 0.0170^{*}$, $W = 0.212$). Post-hoc comparisons did not reveal any pair-wise differences across conditions.

4.4 User Rankings

At the end of the study, participants ranked the visualizations from 1-4 (with 1 being the best). Both the *searcher* and *sorter* were asked to provide these rankings based on their overall experience. It is important to note that while the *sorter* had exposure to different visualizations, the *searcher* did not. Instead, they viewed a panel that displayed how their interfaces were presented to the *sorter* (see Section 3.2.4). This arrangement enabled them to provide their rankings for the different

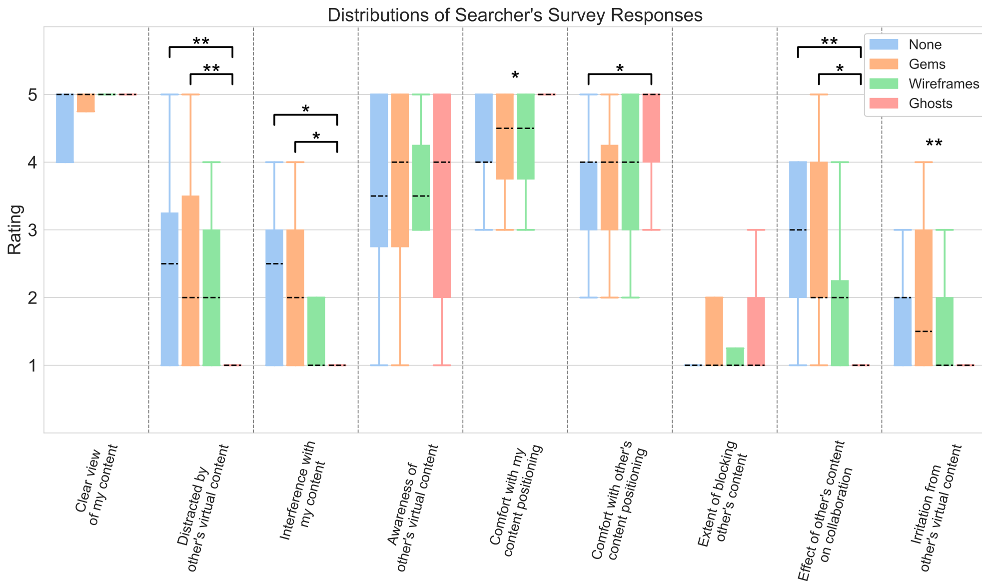

Fig. 7. Boxplots for *searcher's* survey responses

conditions. Figure 9a displays the rankings for visualization, as well as the number of *searchers* and *sorters* that contributed to the individual rankings in parenthesis.

A Friedman test yielded a main effect between VISUALIZATIONS and their rankings ($\chi^2(3) = 52.088$, $p < 0.001^{***}$) and a high consensus among users ($W = 0.543$). *Ghosts* emerged as the top choice for 24 of the 32 participants, while *wireframes* was the favorite for 5 participants and second choice for 21 participants. Interestingly, *gems* received a lower ranking than *none*, which was favored by three participants. Post-hoc analysis revealed *ghosts* was ranked significantly higher compared to *none* ($Z = 1.875$, $p < 0.001^{***}$) and *gems* ($Z = 2.031$, $p < 0.001^{***}$); similarly *wireframes* was ranked higher compared to *none* ($Z = 1.031$, $p = 0.0083^{**}$) and *gems* ($Z = 1.188$, $p = 0.0014^{**}$).

Given the asymmetric design of our experiment, we further divided the rankings based on user roles: *searcher* and *sorter*. Figure 9b displays the ranking distributions for these two roles. A Friedman test showed a main effect between VISUALIZATIONS and their rankings for the *searcher* ($\chi^2(3) = 31.725$, $p < 0.001^{***}$, $W = 0.661$) and the *sorter* ($\chi^2(3) = 22.050$, $p < 0.001^{***}$, $W = 0.459$). The results follow a somewhat similar trend as the overall rankings. *Ghosts* were ranked significantly higher than *none* by both the *searcher* ($Z = 2.250$, $p < 0.001^{***}$) and the *sorter* ($Z = 1.500, p = 0.0061 **$). *Ghosts* was also ranked significantly higher than *gems* by both the *searcher* ($Z = 2.063$, $p < 0.001^{***}$) and the *sorter* ($Z = 1.500, p < 0.001^{***}$). While the *searcher* ranked *wireframes* significantly higher than *none* ($Z = 1.313$, $p = 0.0242^{*}$), this was not observed for the *sorter*. The *sorter* ranked *wireframes* significantly higher than *gems* ($Z = 1.250$, $p = 0.0370^{*}$).

A Spearman's rank correlation conducted for all experimental trials (4 x 16) revealed a correlation of $\rho = 0.55$ ($p < 0.001^{*}$) between the amount of occlusion and the *searcher's* rankings. This denotes a moderate positive correlation, suggesting that as occlusions intensified, *searchers* were more likely to rank visualizations lower (indicating a worse experience). Despite not being instructed to base their rankings on specific criteria such as occlusion, it appears that occlusion significantly impacted the *searchers'* assessments. Section 4.5 contains comments that provide more insights into this observation. Interestingly, *sorters* on average favored the *none* condition over the *gems*

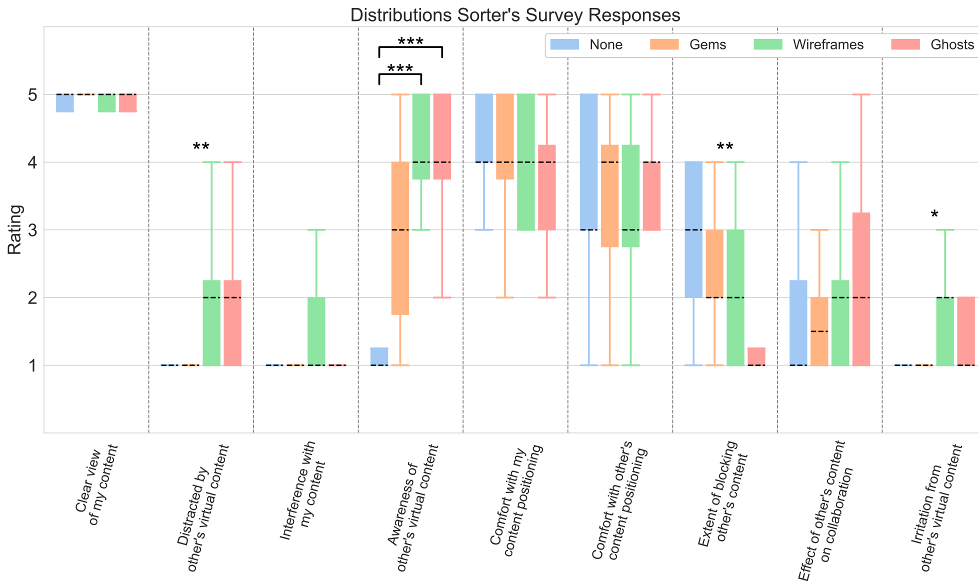

Fig. 8. Boxplots for *sorter's* survey responses

condition. We conducted Wilcoxon signed-rank tests for each condition pair across different user roles, such as comparing *none* from the *searcher* to *none* from the *sorter*. Our findings revealed no main effect, indicating that differences in ranking preferences between roles might be due to random variation, rather than systematic.

4.5 Participants' Comments

At the end of the study, we conducted semi-structured interviews during which participants were asked to explain their rankings of the various VISUALIZATIONS. For *searchers*, a clear, unobstructed view was paramount, leading them to rank the *none* and *gems* conditions lower than the *wireframes* and *ghosts* options. Conversely, *sorters* made their decisions considering a balance between visual disruption and their awareness of the *searcher's* content.

We conducted a thematic analysis to identify themes for each role. The thematic analysis was conducted by a single coder and included a total of 16 transcripts, each corresponding to one study session. The total duration of these sessions was 88 minutes (i.e., total interview time). Initially, participant comments were categorized based on the specific visualization conditions and their corresponding rankings. This was followed by a comprehensive review of the comments, during which initial codes were assigned to them summarizing key ideas. These codes were then grouped and analyzed to form a preliminary set of themes. Finally, we revisited the data set to ensure that these themes were representative of the coded extracts and the data as a whole. This step involved refining the themes, and where necessary, splitting, combining, or discarding those with less significance in the data. In the subsequent sections, we provide representative comments of participants to explain the rationales behind the rankings of each role.

4.5.1 Searcher's Preferences. Searchers generally valued unobstructed views, and noted perceptual challenges with *gems* and *wireframes*.

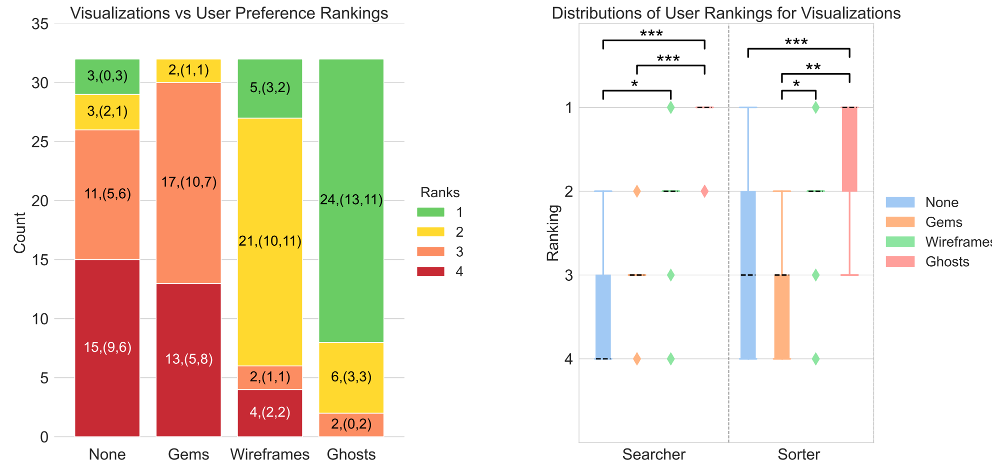

(a) User rankings for visualizations. Values in parenthesis indicate the number of (*searcher*, *sorter*) rankings.

(b) User rankings for visualizations per role

Fig. 9. Results of participants ranking the visualizations from (1) highest to 4 (lowest).

Importance of Unobstructed View. *Searchers* expressed discomfort and decreased satisfaction when their view was blocked by cubes in the *none* condition (*"In this one [none] the cubes just blocked my view a lot."*, P1-*searcher*) and (*"This one [none] was last because when the other person does not see where your panels are it just completely blocks it [the view]."*, P14-*searcher*). *Searchers* preferred visualizations like the *ghosts* that resulted in less occlusion. When contrasted against *none* and *gems*, they found them to cause fewer occlusions (*"I'd say the last one [ghosts] was my favorite because he could see the full window."*, P2-*searcher*) and (*So the first [ghosts] definitely he was very mindful of where my boxes were, the windows were. So he was not putting anything on [them]."*, P5-*searcher*).

Deception of Gems. When using *gems* as visual indicators, *searchers* frequently felt the need to exert additional effort to ensure that *sorters* avoided occlusions, especially when compared to the *none* condition (*"I felt like the diamonds [gems] would give you more harm than good. So he [the sorter] would put it somewhere near the diamond and ask whether it was blocking or not. It adds like extra time and effort."*, P4-*searcher*) and (*"the fourth [ranked] one [gems], she actually placed a lot [of cubes] blocking the charts [panels], though, I think she put it in front of the diamond in her perspective. So the diamond is behind the cube, which actually eventually blocks my view."*, P6-*searcher*).

Depth Perception with Wireframes. While a majority of *searchers* ranked *wireframes* as their second most preferred visualization, some disliked it (*"Even though he had a better idea of where the windows were, he kept putting them in front of them still. Like he couldn't tell he was pushing too far forward."*, P12-*searcher*) and (*"He was constantly lining the boxes over the windows [panels]."*, P5-*searcher*).

4.5.2 Sorter's Preferences. *Sorter's* acknowledged the need for visual cues and boundaries for achieving unobstructed layouts. They appreciated the guidance that *ghosts* and *wireframes* provided. Some *sorter's* preferred the *none* condition for its quality of no visual disruption.

Valuing Clarity over Occlusions. The absence of any visual aids in the *none* category was sometimes preferable to *sorters*, as it provided a clear view without obstructions (*"It was clear to see"*, P7-*sorter*) and (*"didn't think about it [panels] at all and that was good"*, P13-*sorter*).

Insufficient Visual Cue. When it came to using *gems* as visual indicators, many *sorters* found these indicators not to be particularly informative in making decisions about where to place objects. They frequently mentioned the challenges they faced in discerning the exact positions, sizes, or boundaries of the visual elements. This led to uncertainty about where objects should be placed (*"The gems were the worst, like I would try to avoid the gems but still be in the region of where the actual information [panels] was without realizing it."*, P12-*sorter*).

Perceptual Challenges with Wireframes. While considering the *wireframes* visualization, similar to the *searchers*, *sorters* expressed challenges discerning if objects were inside or outside the frame. This suggests issues with depth perception and boundary clarity. (*"It was harder to tell if it [cube] was inside the thing [panel]."*, P2-*sorter*) and (*"It is the same as this one [ghosts] but sometimes I couldn't see the borders"*, P9-*sorter*). This finding aligns with the results of Livingston et al. [31] where the *wireframe* condition was less effective at conveying the depth of the virtual object.

Preference for Holistic Visual Indicators. Similar to the *searchers*, *sorters* indicated a preference for visualizations that provided a complete view, such as *ghosts* and *wireframes*. This method allowed them to accurately judge where to place items without obstructing the other party's view (*"Seeing the full thing is just helpful, I could actually tell if it was intersecting the plane."*, P2-*sorter*) and (*"These ones [wireframes + ghosts] allowed me to understand where I do not have to put the content, so these ones were like easier for me"*, P11-*sorter*).

5 Discussion

Our work explores how to jointly display personal and shared interfaces for collaborative MR scenarios, leveraging visualization that reveal different levels of information about others' personal interfaces to avoid *virtual-virtual conflicts*. Participants prefer visualizations that indicate the position, size, and shape of others' virtual interfaces.

5.1 Reflection on Results

In the absence of any visualization (i.e., the *none* condition), most occlusions occurred. Users, in general, did not prefer this setting. While at first glance, this seemed like a natural result, we were surprised to find that this was also true for the *sorter*, who was actually the person not experiencing occlusions. In multi-user MR collaborative environments, especially those involving tasks where each participant has partial awareness, this means that it is beneficial for the group as a whole to reveal some level of information about personal content to others.

Participants generally preferred the *ghosts* visualization. This visualization not only conveys the position, size, and shape of other's virtual elements, but also slightly occludes the scene where the personal virtual content is, in contrast to the *wireframes* visualization. Participants mentioned that the visualization appeared less intrusive due to its semi-transparent look, which is surprising since the *wireframes* conveys similar information, but enables a better view of the surroundings. We believe that with different tasks and types of personal interfaces (such as 3D), participants' preferences might vary. For instance, they might opt for less visually intrusive designs for visualizations of 3D personal interfaces. We hope to investigate this aspect in the future.

Finally, a few participants, particularly *sorters*, preferred the *none* visualization. They found the visualization of others' personal content unnecessary and preferred only handling their own content. This highlights that even though the visualizations are beneficial, they might still be perceived as visually disruptive. We hope to explore this balance between visual disruption and extent of information, in the future.

5.1.1 Visualization Impact or Order Effect. In our study we decided to present the VISUALIZATIONS in a predetermined progressive order, starting from VISUALIZATIONS that revealed less details about the personal interfaces (i.e., high privacy preserving), and moving to VISUALIZATIONS that exposed more (i.e., low privacy preserving). The rationale for this decision is detailed in Section 3.2.5.

While this design could have caused order effects, feedback from participants (as detailed in our thematic analysis in subsection 4.5), and our strategy to mitigate order effects using varied layouts (as discussed in section 3.2.5), lead us to conclude that the impact observed was predominantly due to the VISUALIZATIONS themselves rather than their presentation order. For example, *wireframe* (the third condition) was perceived as a significant improvement over *gems* (the second condition) because it revealed more information and enabled *sorters* to place cubes in a much more targeted manner. In the subjective ratings, for example, we did *not* see a strong difference between *gems* and *none* because *sorters* felt "deceived" by the fact that *gems* did not indicate size.

5.2 Seeing Is Believing

Although *sorters* were informed that *searchers* had panels in front of them, they seemed to completely overlook this in the *none* condition. This is supported by the high prevalence of occlusions in the *none* condition. We hypothesize that this happened because users often treat virtual objects as physical ones, only acknowledging their existence if visually apparent. This observation aligns with findings from Lebeck et al. [27]. We believe that this finding particularly highlights the need for visualizations that indicate personal objects. In a conversation with one *sorter* (P15-*sorter*), who ranked *none* as the first choice, the following exchange took place:

> ***Sorter***: *"This [none] is one, since there are no objects"*
> ***Experimenter***: *"But, you don't care about occlusions for his [the searcher's] panels?"*
> ***Sorter***: *"What kind of panels? There were no panels."*
> ***Experimenter***: *"But he still had those panels?"*
> ***Sorter***: *"I don't care if I don't see."*

5.3 Design Implications

In the following, we distill a set of design implications based on the results of our study.

5.3.1 Inform Users about Position and Shape. Our findings indicate that merely indicating the existence of personal objects is insufficient to avoid occlusion and distraction. Future systems should visualize the position *and* shape of personal objects. Furthermore, indicating where relevant content (e.g., text) for others exists, for example by incorporating semi-transparent elements, is preferred by users.

5.3.2 Need for Constraints. Even the most informative visualizations, such as *ghosts*, led to occlusions. Furthermore, in some situations, participants just disregarded the visualizations (*"I didn't think about them [the gems] at all, but I think I should have*, P13-*sorter*). This suggests that visualizations alone might not suffice. Systems that have both personal and shared interfaces should account for this, e.g., by introducing constraints that prevent overlap and occlusion. For instance, the introduction of spatial constraints can prevent users from placing virtual interfaces in locations where occlusions would occur. Specifically, these constraints can function like invisible barriers around a users' personal interfaces, which others cannot intrude upon. Additionally, it is crucial that users, whose interfaces are at risk of being occluded, have the authority to set and adjust these constraints to ensure clear visibility and usability of their virtual elements.

5.3.3 No One-Fits-All Solution. Participants displayed diverse preferences between the *wireframes* and *ghosts* conditions, highlighting the importance of individual tastes in choosing visualizations. In the semi-structured interviews, participants were prompted to propose a new visualization. Ideas varied from altering the cube's color to signal occlusion, to triggering visualizations only when an occlusion was imminent. It's conceivable that users might prefer a combination of these visualizations. For instance, users might assign different visualizations to different people's personal interfaces. We believe that future systems should enable some degree of personalization for users to tailor to this preference.

5.3.4 Taking Each Other's Perspective: Privacy-Aware Casting. Users often do not recognize when they are inadvertently blocking the view of others. In the *wireframes* and *ghosts* conditions, for example, when users placed cubes between the panels, many believed they were not causing any occlusion, even though they were. This phenomenon underscores the inherent challenge in perceiving spatial relationships from only one's viewpoint. As one user noted, (*"Even though he had a better idea of where the windows [panels] were, he kept putting them [cubes] in front of them [panels] still. Like he couldn't tell he was pushing too far forward."*, P12-*searcher*).

Our study also revealed that users attempted to confirm the presence of occlusions by placing blocks against panels and verbally checking with their partners. This method, although effective, indicates an inefficient workaround due to the lack of spatial awareness. To address these issues, we propose that future systems could provide users with a view of the other user's perspective when conflicts in spatial arrangement are detected. This feature could be similar to the casting functionality available in many commercial headsets but would require careful implementation to respect privacy. For instance, such systems should avoid showing sensitive content from other users' perspectives and focus solely on sharing spatial configurations of virtual interfaces.

Implementing this feature would involve significant design considerations. It would be crucial to ensure that the additional visual information does not overwhelm users. We recommend activating this perspective-sharing feature only in contexts where occlusion is likely or has been detected, and simplifying the visual information to highlight only critical elements. This targeted approach would help manage cognitive load and streamline the collaborative process.

5.4 Limitations and Future Work

While we believe our work is an important step towards understanding MR environments that contain personal and shared interfaces, there are several challenges that remain. Firstly, we chose a controlled experiment over in-the-wild studies to create comparable conditions. Therefore, our study does not reflect the intricacies of dynamic, real-world MR situations, like workplace settings, for example. We suggest future studies be conducted in actual environments and be longitudinal in nature. Similarly, future studies should investigate the impact of personal interfaces in situations involving more than two users, and different user formations such as side-by-side.

Our current study utilized 2D personal interfaces, which is the standard for many current productivity MR layouts [8, 29, 33]. We hope to expand the visualizations to 3D interfaces in future studies as these could have different implications. For instance, the use of large 3D semi-transparent blobs or 3D ghosts might be less preferred due to increased disruption. This hypothesis could be further examined by varying the intensity and opacity of these 3D ghosts [31].

Additionally, we aim to investigate more privacy-centric methods in subsequent studies. These might include visualizations which do not reveal the number of personal interfaces but rather display a boundary representing the convex hull of the positions of the user's personal interfaces.

Another avenue for future research involves exploring visualizations that change the appearance of the shared objects when they occlude personal objects.

Our study primarily focused on occlusion issues, but future research should also assess the perceived effort and outcomes of user collaboration with these visualizations. Furthermore, studies could examine tasks with differing collaboration requirements. It is possible that users might be indifferent to the occlusion of less significant objects. This could result in different preferences for the visualizations among both user roles.

Furthermore, while our study was conducted exclusively with virtual elements, future works could look into incorporating physical artifacts such as large displays or a whiteboard alongside MR interfaces. Lastly, future studies could delve into multi-modal forms of these visualizations. For instance, systems could verbally notify users of occlusions or provide suggestions for adjustments. Haptic feedback mechanisms, like controller vibrations indicating occlusion, could also be integrated.

Finally, we want to reiterate that we used a predetermined sequence for presenting the visualizations to users in our study. This decision was informed by the nature of the visualizations themselves and our preliminary observations. Our qualitative feedback backs up this decision. However, while our study was carefully designed to minimize order effects (e.g., by varying the layouts), the possibility of their impact cannot be entirely dismissed. Therefore, future research could further explore the implications of randomized sequencing approaches to validate our findings and expand upon our understanding of visualization preferences and their effects.

6 Conclusion

We investigated how different visualizations impact occlusion during collaborative tasks when personal and shared interfaces are present. We conducted a study with 16 dyads performing a collaborative sorting task. Through quantitative and qualitative analyses, we found the fewest occlusions occurred with more informative visualizations. Specifically, the no-visualization (*none*) condition had the most occlusions while the *wireframes* and *ghosts* conditions had significantly fewer. The *ghosts* condition was the most preferred visualization as it revealed complete spatial information about personal objects, avoiding issues with depth perception in the *wireframes* condition. While participants' perceptions and rankings generally aligned with these results, some individuals preferred no visualization, *none*, due to reduced visual clutter. Further, the *gems* visualization, originally intended to reduce occlusions, sometimes led to exacerbating challenges as participants did not have the full information to judge the extent of others' virtual contents. We believe that our findings and the resulting design implications highlight that personal interfaces do not only need to be seen by the user who owns them, but are relevant for the collaborative dynamics as a whole.

References

[1] Huidong Bai, Prasanth Sasikumar, Jing Yang, and Mark Billinghurst. 2020. A user study on mixed reality remote collaboration with eye gaze and hand gesture sharing. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13. https://doi.org/10.1145/3313831.3376550

[2] Blaine Bell, Steven Feiner, and Tobias Höllerer. 2001. View management for virtual and augmented reality. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. 101–110. https://doi.org/10.1145/502348.502363

[3] Mark Billinghurst and Hirokazu Kato. 1999. Collaborative mixed reality. In *Proceedings of the first international symposium on mixed reality*. 261–284.

[4] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. 2019. Cross-device taxonomy: Survey, opportunities and challenges of interactions spanning across multiple devices. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–28. https://doi.org/10.1145/3290605.3300792

[5] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, trends, and outliers: How immersive technologies can facilitate the collaborative analysis of multidimensional data. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12. https://doi.org/10.1145/3173574.3173664

[6] Salvatore Cafiso, Alessandro Di Graziano, and Giuseppina Pappalardo. 2013. Using the Delphi method to evaluate opinions of public transport managers on bus safety. *Safety science* 57 (2013), 254–263. https://doi.org/10.1016/j.ssci.2013.03.001

[7] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 2019. Dataspace: A reconfigurable hybrid reality environment for collaborative information analysis. In *2019 IEEE conference on virtual reality and 3D user interfaces (VR)*. IEEE, 145–153. https://doi.org/10.1109/VR.2019.8797733

[8] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. Semanticadapt: Optimization-based adaptation of mixed reality layouts leveraging virtual-physical semantic connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 282–297. https://doi.org/10.1145/3472749.3474750

[9] Yi Fei Cheng, Hang Yin, Yukang Yan, Jan Gugenheimer, and David Lindlbauer. 2022. Towards understanding diminished reality. In *CHI Conference on Human Factors in Computing Systems* (New Orleans LA USA). ACM, New York, NY, USA, 1–16. https://doi.org/10.1145/3491102.3517452

[10] Hyunsung Cho, Yukang Yan, Kashyap Todi, Mark Parent, Missie Smith, Tanya R. Jonker, Hrvoje Benko, and David Lindlbauer. 2024. MineXR: Mining Personalized Extended Reality Interfaces. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 609, 17 pages. https://doi.org/10.1145/3613904.3642394

[11] Shakiba Davari, Feiyu Lu, and Doug A Bowman. 2022. Validating the benefits of glanceable and context-aware augmented reality for everyday information access tasks. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 436–444. https://doi.org/10.1109/VR51125.2022.00063

[12] Barrett Ens, Joel Lanir, Anthony Tang, Scott Bateman, Gun Lee, Thammathip Piumsomboon, and Mark Billinghurst. 2019. Revisiting collaboration through mixed reality: The evolution of groupware. *International Journal of Human-Computer Studies* 131 (2019), 81–98. https://doi.org/10.1016/j.ijhcs.2019.05.011

[13] Michele Fiorentino, Raffaele de Amicis, Giuseppe Monno, and Andre Stork. 2002. Spacedesign: A mixed reality workspace for aesthetic industrial design. In *Proceedings. International Symposium on Mixed and Augmented Reality*. IEEE, 86–318. https://doi.org/10.1109/ISMAR.2002.1115077

[14] Susan R Fussell, Robert E Kraut, and Jane Siegel. 2000. Coordination of communication: Effects of shared visual context on collaborative work. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*. 21–30. https://doi.org/10.1145/358916.358947

[15] Danilo Gasques, Janet G Johnson, Tommy Sharkey, Yuanyuan Feng, Ru Wang, Zhuoqun Robin Xu, Enrique Zavala, Yifei Zhang, Wanze Xie, Xinming Zhang, et al. 2021. ARTEMIS: A collaborative mixed-reality system for immersive surgical telementoring. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/10.1145/3411764.3445576

[16] Darren Gergle, Robert E Kraut, and Susan R Fussell. 2013. Using visual information for grounding and awareness in collaborative tasks. *Human–Computer Interaction* 28, 1 (2013), 1–39.

[17] Edward T Hall. 1969. The Hidden Dimension: man's use of space in public and private The Bodley Head. *London, Sydney, Toronto* 121 (1969).

[18] Heiko Hecht, Robin Welsch, Jana Viehoff, and Matthew R Longo. 2019. The shape of personal space. *Acta psychologica* 193 (2019), 113–122. https://doi.org/10.1016/j.actpsy.2018.12.009

[19] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2021. Stream: Exploring the combination of spatially-aware tablets with augmented reality head-mounted displays for immersive analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/10.1145/3411764.3445298

[20] Bret Jackson, Linda Lor, and Brianna C Heggeseth. 2024. Workspace Guardian: Investigating Awareness of Personal Workspace Between Co-Located Augmented Reality Users. *IEEE Transactions on Visualization and Computer Graphics* (2024). https://doi.org/10.1109/TVCG.2024.3372073

[21] Allison Jing, Kieran May, Gun Lee, and Mark Billinghurst. 2021. Eye see what you see: Exploring how bi-directional augmented reality gaze visualisation influences co-located symmetric collaboration. *Frontiers in Virtual Reality* 2 (2021), 697367. https://doi.org/10.3389/frvir.2021.697367

[22] Hannes Kaufmann. 2003. Collaborative augmented reality in education. *Institute of Software Technology and Interactive Systems, Vienna University of Technology* (2003), 2–4.

[23] Talha Khan, Edward G Andrews, Paul A Gardner, Arka N Mallela, Jeffrey R Head, Joseph C Maroon, Georgios A Zenonos, Dmitriy Babichenko, and Jacob T Biehl. 2022. AR in the OR: exploring use of augmented reality to support endoscopic surgery. In *ACM International Conference on Interactive Media Experiences*. 267–270. https://doi.org/10.1145/3505284.3532970

[24] Talha Khan, Toby S Zhu, Thomas Downes, Lucille Cheng, Nicolás M Kass, Edward G Andrews, and Jacob T Biehl. 2023. Understanding effects of visual feedback delay in ar on fine motor surgical tasks. *IEEE Transactions on Visualization and Computer Graphics* (2023). https://doi.org/10.1109/TVCG.2023.3320214

[25] Kangsoo Kim, Mark Billinghurst, Gerd Bruder, Henry Been-Lirn Duh, and Gregory F Welch. 2018. Revisiting Trends in Augmented Reality Research: A Review of the 2nd Decade of ISMAR (2008-2017). *IEEE transactions on visualization and computer graphics* 24, 11 (Nov. 2018), 2947–2962. https://doi.org/10.1109/TVCG.2018.2868591

[26] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachselt. 2021. Marvis: Combining mobile devices and augmented reality for visual data analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17. https://doi.org/10.1145/3411764.3445593

[27] Kiron Lebeck, Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2018. Towards security and privacy for multi-user augmented reality: Foundations with end users. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 392–408. https://doi.org/10.1109/SP.2018.00051

[28] Zhen Li, Michelle Annett, Ken Hinckley, Karan Singh, and Daniel Wigdor. 2019. Holodoc: Enabling mixed reality workspaces that harness physical and digital content. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/10.1145/3290605.3300917

[29] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-aware online adaptation of mixed reality interfaces. In *Proceedings of the 32nd annual ACM symposium on user interface software and technology*. 147–160. https://doi.org/10.1145/3332165.3347945

[30] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2020. Design and evaluation of interactive small multiples data visualisation in immersive spaces. In *2020 IEEE conference on virtual reality and 3D user interfaces (VR)*. IEEE, 588–597. https://doi.org/10.1109/VR46266.2020.00081

[31] Mark A Livingston, J Edward Swan, Joseph L Gabbard, Tobias H Hollerer, Deborah Hix, Simon J Julier, Yohan Baillot, and Dennis Brown. 2003. Resolving multiple occluded layers in augmented reality. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.* IEEE, 56–65. https://doi.org/10.1109/ISMAR.2003.1240688

[32] Joan Llobera, Bernhard Spanlang, Giulio Ruffini, and Mel Slater. 2010. Proxemics with multiple dynamic characters in an immersive virtual environment. *ACM Transactions on Applied Perception (TAP)* 8, 1 (2010), 1–12. https://doi.org/10.1145/1857893.1857896

[33] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachselt. 2022. Where should we put it? layout and placement strategies of documents in augmented reality for collaborative sensemaking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–16. https://doi.org/10.1145/3491102.3501946

[34] Tahir Mahmood, Willis Fulmer, Neelesh Mungoli, Jian Huang, and Aidong Lu. 2019. Improving information sharing and collaborative analysis for remote geospatial visualization using mixed reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 236–247. https://doi.org/10.1109/ISMAR.2019.00021

[35] Nicolai Marquardt, Ken Hinckley, and Saul Greenberg. 2012. Cross-device interaction via micro-mobility and f-formations. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. 13–22. https://doi.org/10.1145/2380116.2380121

[36] Paul Marshall, Yvonne Rogers, and Nadia Pantidi. 2011. Using F-formations to analyse spatial patterns of interaction in physical environments. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. 445–454. https://doi.org/10.1145/1958824.1958893

[37] Jens Müller, Roman Rädle, and Harald Reiterer. 2016. Virtual objects as spatial cues in collaborative mixed reality environments: How they shape communication behavior and user task load. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 1245–1249. https://doi.org/10.1145/2858036.2858043

[38] Nels Numan and Anthony Steed. 2022. Exploring user behaviour in asymmetric collaborative mixed reality. In *Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology*. 1–11. https://doi.org/10.1145/3562939.3565630

[39] Thammathip Piumsomboon, Gun A Lee, Jonathon D Hart, Barrett Ens, Robert W Lindeman, Bruce H Thomas, and Mark Billinghurst. 2018. Mini-me: An adaptive avatar for mixed reality remote collaboration. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13. https://doi.org/10.1145/3173574.3173620

[40] Lev Poretski, Ofer Arazy, Joel Lanir, and Oded Nov. 2021. Who owns what? Psychological ownership in shared augmented reality. *International Journal of Human-Computer Studies* 150 (2021), 102611. https://doi.org/10.1016/j.ijhcs.2021.102611

[41] Lev Poretski, Joel Lanir, and Ofer Arazy. 2018. Normative tensions in shared augmented reality. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–22. https://doi.org/10.1145/3274411

[42] Iulian Radu, Tugce Joy, and Bertrand Schneider. 2021. Virtual makerspaces: merging AR/VR/MR to enable remote collaborations in physical maker activities. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*. 1–5. https://doi.org/10.1145/3411763.3451561

[43] Iulian Radu and Bertrand Schneider. 2022. How Augmented Reality (AR) Can Help and Hinder Collaborative Learning: A Study of AR in Electromagnetism Education. *IEEE transactions on visualization and computer graphics* (2022). https://doi.org/10.1109/TVCG.2022.3169980

[44] Shwetha Rajaram, Chen Chen, Franziska Roesner, and Michael Nebeling. 2023. Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17. https://doi.org/10.1145/3544548.3581089

[45] Patrick Reipschlager, Tamara Flemisch, and Raimund Dachselt. 2020. Personal augmented reality for information visualization on large interactive displays. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1182–1192. https://doi.org/10.1109/TVCG.2020.3030460

[46] Jun Rekimoto. 1996. Transvision: A hand-held augmented reality system for collaborative design. In *Proceeding of Virtual Systems and Multimedia*, Vol. 96. 18–20.

[47] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2019. Secure {Multi-User} content sharing for augmented reality applications. In *28th USENIX Security Symposium (USENIX Security 19)*. 141–158.

[48] Mickael Sereno, Xiyao Wang, Lonni Besançon, Michael J Mcguffin, and Tobias Isenberg. 2020. Collaborative work in augmented reality: A survey. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2020), 2530–2549. https://doi.org/10.1109/TVCG.2020.3032761

[49] Kihoon Son, Hwiwon Chun, Sojin Park, and Kyung Hoon Hyun. 2020. C-Space: an interactive prototyping platform for collaborative spatial design exploration. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13. https://doi.org/10.1145/3313831.3376452

[50] Ada V Taylor, Ayaka Matsumoto, Elizabeth J Carter, Alexander Plopski, and Henny Admoni. 2020. Diminished reality for close quarters robotic telemanipulation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 11531–11538. https://doi.org/10.1109/IROS45743.2020.9341536

[51] Laurie M Wilcox, Robert S Allison, Samuel Elfassy, and Cynthia Grelik. 2006. Personal space in virtual reality. *ACM Transactions on Applied Perception (TAP)* 3, 4 (2006), 412–428. https://doi.org/10.1145/1190036.1190041

[52] Stefanie Zollmann, Tobias Langlotz, Raphael Grasset, Wei Hong Lo, Shohei Mori, and Holger Regenbrecht. 2021. Visualization Techniques in Augmented Reality: A Taxonomy, Methods and Patterns. *IEEE transactions on visualization and computer graphics* 27, 9 (Sept. 2021), 3808–3825. https://doi.org/10.1109/TVCG.2020.2986247

Received 2024-02-22; accepted 2024-05-30